

What is Claimed

1. A method for enabling the execution of at least an I/O operation while providing a snapshot copy of a storage system, said method comprises the steps of:
 - a) performing on-line at least a primary task of said I/O operation, wherein said primary task is performed using a journal;
 - b) generating a response message ending the execution of said I/O operation; and,
 - c) performing off-line secondary tasks of said I/O operation.
2. The method of claim 1, wherein said storage system includes at least: a host, a journal, a snapshot storage element, a production storage element.
3. The method of claim 2, wherein said snapshot storage element is at least one of: a virtual volume, a physical storage device.
4. The method of claim 2, wherein said production storage element is at least one of: a virtual volume, a physical storage device.
5. The method of claim 1, wherein said journal includes at least one non-volatile random access memory (NVRAM) unit.

6. The method of claim 1, wherein said I/O operation is a write request initiated by the host.
7. The method of claim 6, wherein said performing on-line at least a primary task further comprises the steps of:
 - a) writing a data chunk included in said write request into said journal;and,
 - b) saving a destination address designated in said write request in a changes table.
8. The method of claim 6, wherein said generating a response message further comprises the step of:
 - sending said response message to said host.
9. The method of claim 6, wherein said performing off-line secondary tasks further comprises the steps of:
 - a) checking if said data chunk residing in the snapshot storage element was modified since a last time said snapshot copy was created;
 - b) copying said data chunk from a location in the production storage element to said snapshot storage element and further copying said data chunk from said journal to a location in said production storage element, if said data chunk has not been modified; and,

c) copying said data chunk from said journal to said production storage element, if said data chunk has been modified.

10. The method of claim 9, wherein the location in said production storage element is determined by said destination address.

11. The method of claim 10, wherein said destination address is converted to a physical address if said production storage element is a virtual volume.

12. The method of claim 1, wherein said I/O operation is a read request initiated by the host computer.

13. The method of claim 12, wherein said performing on-line at least a primary task further comprises the steps of:

a) checking if a data chunk requested to be read resides in said journal;

and,

b) retrieving said data chunk from said journal and further sending said data chunk to said host, if said data chunk resides in said journal.

14. The method of claim 13, wherein checking if said data chunk resides in said journal further comprises the step of:

checking whether the changes table includes an entry associated with said data chunk.

15. The method of claim 13, wherein said performing off-line secondary tasks further comprises the steps of:

a) retrieving said data chunk from the production storage element, if said data chunk does not reside in said journal; and,

b) sending said data chunk to said host.

16. The method of claim 15, wherein said data chunk is retrieved from a location designated by a source address included in said read request.

17. The method of claim 16, wherein said source address is converted to a physical address if said production storage element is a virtual volume.

18. A computer-readable medium having stored thereon computer executable code enabling the execution of at least an I/O operation while providing a snapshot copy of a storage system, said executable code for performing the steps of:

a) performing on-line at least a primary task of said I/O operation, wherein said primary task is performed using a journal;

b) generating a response message ending the execution of said I/O operation; and,

c) performing off-line secondary tasks of said I/O operation.

19. The computer executable code of claim 18, wherein said storage system comprises at least: a host, a journal, a snapshot storage element, a production storage element.
20. The computer executable code of claim 19, wherein said snapshot storage element is at least one of: a virtual volume, a physical storage device.
21. The computer executable code of claim 19, wherein said production storage element is at least one of: a virtual volume, a physical storage device
22. The computer executable code of claim 18, wherein said journal includes at least one non-volatile random access memory (NVRAM) unit.
23. The computer executable code of claim 18, wherein said I/O operation is a write request initiated by the host.
24. The computer executable code of claim 23, wherein said performing on-line at least a primary task further comprises the steps of:
 - a) writing a data chunk included in said write request into said journal;and,
 - b) saving a destination address designated in said write request in a changes table.

25. The computer executable code of claim 23, wherein said generating a response message further comprises the step of:

 sending said response message to said host.

26. The computer executable code of claim 23, wherein said performing off-line secondary tasks further comprises the steps of:

 a) checking if said data chunk resides in the snapshot storage element was modified since a last time said snapshot copy was created;

 b) copying said data chunk from a location in the production storage element to said snapshot storage element and further copying said data chunk from said journal to a location in said production storage element, if said data chunk has not been modified; and,

 c) copying said data chunk from said journal to said production storage element, if said data chunk has been modified.

27. The computer executable code of claim 26, wherein the location in said production storage element is determined by said destination address.

28. The computer executable code of claim 27, wherein said destination address is converted to a physical address if said production storage element is a virtual volume.

29. The computer executable code of claim 18, wherein said I/O operation is a read request by the host.

30. The computer executable code of claim 29, wherein said performing at least a primary task further comprises the steps of:

a) checking if a data chunk requested to be read resides in said journal;

and,

b) retrieving said data chunk from said journal and further sending said data chunk to said host, if said data chunk resides in said journal.

31. The computer executable code of claim 30, wherein checking if said data chunk resides in said journal further includes:

checking whether the changes table includes an entry associated with said data chunk.

32. The computer executable code of claim 30, wherein said performing off-line secondary tasks further comprises the steps of:

a) retrieving said data chunk from the production storage element, if said data chunk does not reside in said journal; and,

b) sending said data chunk to said host.

33. The computer executable code of claim 32, wherein said data chunk is retrieved from a location designated by a source address included in said read request.
34. The computer executable code of claim 33, wherein said source address is converted to a physical address if said production storage element is a virtual volume.
35. An apparatus for execution of at least an I/O operation with minimal latency while providing a snapshot copy of a storage system, said apparatus comprising:
- means for receiving said at least an I/O operation from a host;
 - means for performing on-line at least one primary task of said at least an I/O operation;
 - means for performing off-line at least a secondary task of said at least an I/O operation;
 - means for controlling a snapshot storage element;
 - means for controlling a production storage element; and,
 - means for controlling a journal.
36. The apparatus of claim 35, wherein said snapshot storage element is at least one of: a virtual volume, a physical storage device.
37. The apparatus of claim 35, wherein said production storage element is at least one of: a virtual volume, a physical storage device.

38. The apparatus of claim 35, wherein said physical storage device comprises at least one of: tape drive, tape library, optical drive, disk, redundant array of independent disks (RAID).
39. The apparatus of claim 35, wherein said journal includes at least one non-volatile random access memory (NVRAM) unit.
40. The apparatus of claim 35, wherein said I/O operations comprise at least one of: write request, read request.
41. The apparatus of claim 35, wherein said primary task is performed on-line using said journal.
42. The apparatus of claim 35, wherein said apparatus is a storage controller.
43. The apparatus of claim 35, wherein said apparatus is a virtualization switch connected in a storage area network (SAN).
44. The apparatus of claim 35, wherein said I/O operation is a read request initiated by said host.

45. The apparatus of claim 44, wherein said performing said primary task further comprises the steps of:
- a) checking if a data chunk requested to be read resides in said journal;
- and,
- b) retrieving said data chunk from said journal and further sending said data chunk to said host, if said data chunk resides in said journal.
46. The apparatus of claim 45, wherein checking if said data chunk resides in said journal further comprises:
- checking whether the changes table includes an entry associated with said data chunk.
47. The apparatus of claim 45, wherein said performing off-line secondary tasks further comprises the steps of:
- a) retrieving said data chunk from the production storage element, if said data chunk does not reside in said journal; and,
 - b) sending said data chunk to said host.
48. The apparatus of claim 35, wherein said I/O operation is a write request initiated by said host.
49. The apparatus of claim 48, wherein said performing on-line at least a primary task further comprises the steps of:

- a) writing a data chunk included in said write request into said journal;,
- b) saving a destination address designated in said write request in a changes table; and,
- c) sending a response message to said host.

50. The apparatus of claim 48, wherein said performing off-line secondary tasks further comprises the steps of:

- a) checking if said data chunk resides in the snapshot storage element was modified since a last time said snapshot copy was created;
- b) copying said data chunk from a location in the production storage element to said snapshot storage element and further copying said data chunk from said journal to a location in said production storage element, if said data chunk has not been modified; and,
- c) copying said data chunk from said journal to said production storage element, if said data chunk has been modified.